High Frequency Trading and Extreme Price Movements*

Jonathan Brogaard, Allen Carrion, Thibaut Moyaert, Ryan Riordan, Andriy Shkilko, Konstantin Sokolov

First version: November 2014
This version: June 2016

Abstract: Are endogenous liquidity providers (ELPs) unreliable in times of market stress? We examine the activity of a common ELP type – high frequency traders (HFTs) – around extreme price movements (EPMs). We find that on average HFTs provide liquidity during EPMs by absorbing imbalances created by non-high frequency traders (nHFTs). Yet HFT liquidity provision is limited to EPMs in single stocks. When several stocks experience simultaneous EPMs, HFTs do not supply liquidity. There is little evidence of HFTs causing EPMs. HFTs earn positive revenues during the average EPM.

Internet appendix: http://bit.ly/28Zbu5C

* We thank Amber Anand, Hendrik Bessembinder, Phelim Boyle, Sabrina Buti, Jean-Edouard Colliard, Emmanuel Gobet, Kingsley Fong, Nathan Halmrast, Terrence Hendershott, Thierry Foucault, Katya Malinova, Olena Nikolsko-Rzhevska, Maureen O'Hara, Michael Pagano, Andreas Park, Roberto Pascual, Fabricio Perez, Wing Wah Tham, Jun Uno, Kumar Venkataraman, Haoxiang Zhu, and conference participants at the AFA, Erasmus Liquidity Conference, FMA, the Financial Risks International Forum, FMA Europe, MARC, and NFA for insightful comments. We are grateful to Nasdaq OMX for providing the data. Part of the research presented in this paper was performed while Allen Carrion served as a Visiting Financial Economist at the U.S. Securities and Exchange Commission. The Securities and Exchange Commission, as a matter of policy, disclaims responsibility for any private publication or statement by any of its employees. The views expressed herein are those of the author and do not necessarily reflect the views of the Commission or of the author's colleagues on the staff of the Commission. Moyaert acknowledges financial support from Actions de Recherche Concertées.

Jonathan Brogaard, University of Washington, brogaard@uw.edu; Allen Carrion, University of Utah and AIQ, LLC, al.carrion@business.utah.edu; Thibaut Moyaert, Louvain School of Management, thibaut.moyaert@uclouvain.be; Ryan Riordan, Queen's University, ryan.riordan@queensu.ca; Andriy Shkilko, Wilfrid Laurier University, ashkilko@wlu.ca; Konstantin Sokolov, Wilfrid Laurier University, soko7760@mylaurier.ca.

1. Introduction

Since the 2008-09 financial crisis, the fragility of financial markets has been widely debated. The May 2010 Flash Crash highlighted one aspect of this debate: the relation between extreme price movements (EPMs) and certain forms of electronic trading, namely high frequency trading (HFT). EPMs (or price jumps) have long been a topic of study in the finance literature, with a number of papers suggesting that they may have adverse effects on markets. For instance, EPMs may impair risk management (Duffie and Pan, 2001), derivative pricing (Eraker, Johannes and Polson, 2003) and portfolio allocation (Jarrow and Rosenfield, 1984; Liu, Longstaff and Pan, 2003). Given the importance of EPMs and the ubiquity of HFT in today's markets, we examine in detail the relation between EPMs and HFT.

In modern markets high frequency traders (HFTs) play an important, if not the dominant, role in providing liquidity (Hasbrouck and Saar, 2013; Menkveld, 2013; Malinova, Park and Riordan, 2014, Conrad, Wahal and Xiang, 2015). Generally, the rise of HFT has been accompanied by a reduction in trading costs (Angel, Harris and Spatt, 2011; Jones, 2013; Harris, 2013) and an increase in price efficiency (Carrion, 2013; Brogaard, Hendershott and Riordan, 2014; Chaboud, Chiquoine, Hjalmarsson and Vega, 2014). Yet HFTs are endogenous liquidity providers (ELPs), typically without an obligation to stabilize markets during stressful periods. Do HFTs provide liquidity selectively and only during periods of market calm?

Our main finding is that, on average, HFTs trade in the opposite direction of rapidly developing extreme price movements and supply liquidity to non-high frequency traders (nHFTs) by absorbing their trade imbalances. This result holds even during the largest EPMs and during the times when nHFTs demand substantial amounts of liquidity. Notably, HFTs supply liquidity both to the EPMs that eventually reverse and the EPMs that result in permanent price changes. An average HFT trade during extreme price movements provides liquidity to aggressive, occasionally informed nHFTs. HFT liquidity demand also increases during EPMs, but the increase in liquidity supply is of greater magnitude, resulting in HFTs supplying net liquidity during EPMs.

While we find that HFTs provide liquidity during EPMs it is possible that HFTs also trigger them. Chordia et al. (2013) write: "There is growing unease on the part of some market observers that […] violent price moves are occurring more often in financial instruments in which HFTs are active." Golub, Keane and Poon (2013) report that mini-crashes in individual stocks have increased in recent years and suggest a link between these crashes and HFT. Leal, Napoletano, Roventini and Fagiolo (2014) model a market in which HFTs play a fundamental role in generating flash crashes. Media reports and industry commentary also often draw a causal link between HFT and EPMs. In October 2015 Timothy Massad, chair of the U.S. Commodities Futures Trading Commission (CFTC), expressed a concern over sudden large price movements and linked them to high-speed computerized trading.[1] We examine this link using a probit analysis of EPM determinants. We find no evidence of HFTs triggering EPMs.

On average HFTs stabilize prices during EPMs. Are there limits to this stabilizing behavior? Theory and empirical research suggests that ELPs may stop making markets during stressful periods when liquidity is most needed (Raman, Robe and Yadav, 2014; Anand and Venkataraman, 2015; Bongaerts and Van Achter, 2015; Cespa and Vives, 2015; Korajczyk and Murphy, 2015). We too find limits to HFT liquidity provision during EPMs. Specifically, HFTs stop providing liquidity to nHFTs when more than one stock simultaneously undergoes an EPM (co-EPMs). This result is likely due to capital constraints and reduced cross-asset hedging opportunities inhibiting the HFTs' ability to provide liquidity. Focusing on one exceptionally

[1] "US regulator signals bid to curb high-speed trading," by Gregory Meyer and Joe Rennison, Financial Times, October 21, 2015.

large co-EPM – the 2010 Flash Crash – Kirilenko, Kyle, Samadi and Tuzun (2015) also find that HFTs withdrew from liquidity provision.

Even though EPMs occur quickly, they consist of multiple sequential trades. If HFT algorithms had been designed to stop providing liquidity during EPMs, technology would have allowed them to withdraw limit orders as EPMs develop. Yet the results show that the algorithms are designed to remain in the market. To understand what motivates such design, we refer to the literature that examines contrarian liquidity provision. Hendershott and Seasholes (2007), Nagel (2012) and So and Wang (2014) find that providing liquidity against price movements that eventually reverse is a profitable strategy. Notably, many EPMs occur during large intraday price reversals. If the profits derived from such reversals are sufficiently high, HFTs should remain in the market to capture them. Consistent with this intuition, we find that HFT profits on days with EPMs are substantially higher than they are on an average day.

Our analysis generalizes the results of studies that examine the 2010 Flash Crash (e.g., Easley, Lopez de Prado, and O'Hara, 2012; Kirilenko, Kyle, Samadi and Tuzun, 2015; and Menkveld and Yueshen, 2015). We examine more than 45,000 EPMs during a two-year period. The data span 2008 and 2009 and therefore capture the heightened intraday volatility in financial markets during the 2008 financial crisis as well as more normal conditions. Overall, we fail to find evidence that HFTs cause EPMs. We also show that HFTs provide liquidity to nHFTs during an average EPM and make money from the price reversals that follow. This said, HFT liquidity provision is constrained when multiple stocks undergo simultaneous EPMs. Researchers and regulators should account for this characteristic in market design applications.

2. Data, EPM detection and summary statistics

2.1. HFT data

The HFT data come from NASDAQ and span two years: 2008 and 2009. These data have been previously used by Carrion (2013), Brogaard, Hendershott and Riordan (2014), and O'Hara, Yao and Ye (2014), among others. For each trade the dataset contains an indicator for whether an HFT or an nHFT participates on the liquidity-supplying or the liquidity-demanding side of the trade. When preparing the data NASDAQ identified 26 firms that act as independent HFT proprietary trading firms based on its knowledge of the firm's activity. A firm is identified by NASDAQ as an HFT if it trades frequently, holds small intraday inventory positions, and ends the day with a near zero inventory. By 2008, the HFT industry had largely matured, making the results applicable in today's market. An additional benefit to using these data is that HFTs on NASDAQ have no obligation to stabilize prices during stressful times (Bessembinder, Hao and Lemmon, 2011; Clark-Joseph, Ye and Zi, 2016) and so are ideal to study liquidity provision by ELPs.

The data allow us to directly observe HFT liquidity provision and demand. We are subject to the same limitations as the abovementioned studies, mainly that we cannot observe individual HFT activity and that we only observe trading on NASDAQ. Although trades on NASDAQ make up 30-40% of all trading activity in the sample stocks, it is a possible that during EPMs HFTs provide liquidity on NASDAQ while taking it from the other markets. We are unable to refute this possibility. Nonetheless, we believe that such liquidity transfer is unlikely as liquidity provision on NASDAQ is not systematically more attractive than it is on other venues during the sample period.

2.2. EPM identification

We identify EPMs as extreme changes in the National Best Bid and Offer (NBBO) midquotes. The use of midquotes instead of trade prices allows us to reduce the effect of the bid-ask bounce. The results are similar when we use trade prices. We obtain the midquotes from the NYSE Trade and Quote database (TAQ) after adjusting the data according to Holden and Jacobsen's (2014) recommendations. Specifically, we (i) interpolate the times of trades and the times of NBBO quotes within a second, (ii) adjust for withdrawn quotes, (iii) delete locked and crossed NBBO quotes, and (iv) delete trades reported while the NBBO is locked or crossed. To avoid focusing on price dislocations that may be caused by market opening and closing procedures, we only consider trading activity between 9:35 a.m. and 3:55 p.m.

Using the filtered TAQ midquotes, we compute 10-second absolute midquote returns. The choice of the 10-second sampling frequency is based on two offsetting considerations. On the one hand, detecting EPMs that result from brief liquidity dislocations requires a relatively short sampling interval. On the other hand, a sampling interval that is too short may split an EPM into several small price changes that are not large enough to be captured by the identification procedure. The choice of 10-second intervals is a compromise between these two considerations. As a robustness check, the main analyses are repeated for several alternative interval lengths: 1 second, 5 seconds, 30 seconds, and 1 minute. The results are qualitatively similar.

The NASDAQ HFT dataset contains 120 stocks divided into three size categories: large, medium and small, with 40 stocks in each category. Medium and small stocks trade rather infrequently, and there are usually insufficient observations to draw statistically robust conclusions about HFT and nHFT activity during the sampling intervals. The main analysis therefore focuses on the 40 largest stocks. In a similar application, and driven by similar considerations, Andersen, Bollerslev, Diebold and Ebens (2001) also focus on the largest stocks when detecting EPMs. The sample of 40 largest stocks contains over 45.4 million 10-second intervals. In a robustness test discussed in a later section, we use volume bucketing that allows us to examine extreme price movements in medium and small stocks.

We define an EPM as an interval that belongs to the 99.9th percentile of 10-second absolute midpoint returns for each stock. That is, out of 45.4 million 10-second intervals, we identify 45,406 intervals with the largest returns as EPMs. The intuitive nature of the 99.9 technique is appealing, but the technique has two limitations. First, the 99.9 cutoffs are stock-specific and therefore implicitly assume that each stock is equally likely to undergo an EPM. Consequently, the 99.9 technique may (over-) under-sample stocks that are (less) more prone to EPMs. The second limitation is that the technique is agnostic to volatility conditions and therefore tends to oversample periods of high volatility. We suggest that understanding HFT behavior is relevant regardless of whether the EPM is accompanied by high volatility. Nevertheless, to formally address this limitation, we repeat the analysis using another EPM detection technique, the Lee and Mykland’s (2012) methodology, which accounts for contemporaneous volatility. The results obtained using this methodology are reported in the internet appendix and are qualitatively similar to those from the 99.9 technique.

Finally, in unreported results, we find that returns in the 99.9th percentile closely correspond to the 99.9th percentile of trade imbalances. EPM identification that focuses on the largest imbalances rather than the largest returns produces a similar sample.

2.3. Summary statistics

Table 1 reports the descriptive statistics for the sample of 45,406 EPMs in Panel A and, for comparison, the full sample of 10-second intervals in Panel B. The statistics expectedly show that returns, trading activity, and spreads are considerably larger during the EPMs than during an average 10-second period. The average absolute EPM return is 0.484%, which is more than 17 times (or more than 10 standard deviations) larger than the full-sample return. Trading activity is also substantially higher; increasing from 18 trades per 10 seconds to 73 trades. Dollar trading volume increases from $76,285 to $473,232, and share volume increases by a similar magnitude. Finally, the quoted and relative spreads nearly double during EPMs.

The number of positive EPMs is approximately equal to the number of negative EPMs. In untabulated results, we find that EPM characteristics such as the absolute return magnitude, trading volume, and quoted spreads are similar for positive and negative EPMs. HFT and nHFT behavior is also similar. The results reported in the remainder of this manuscript combine positive and negative EPMs.

INSERT TABLE 1 ABOUT HERE

Figures 1 and 2 report the time series EPM distributions. Figure 1 reports the intraday frequency of EPMs, with 53.8% of the events occurring in the first hour of trading. This pattern is consistent with studies that document high price volatility and information uncertainly in the morning hours (Chan, Christie and Schultz, 1995; Egginton, 2014). The remaining EPMs are distributed relatively evenly throughout the day, with a moderate increase near the end of the day.[2] Figure 2 plots the daily frequency of EPMs during the 2008-2009 sample period. Most of the EPMs in the sample (65.1%) occur during the months of September, October and November of 2008, the height of the financial crisis.

INSERT FIGURES 1 AND 2 ABOUT HERE

3. HFT and nHFT activity around EPMs

In this section we show that HFTs provide liquidity to nHFTs during a typical EPM, even when the EPM is very large and even when the price change is permanent. However, HFT liquidity provision is unreliable when several stocks undergo simultaneous EPMs. We also show that liquidity provision during an average EPM is profitable, yet we find no evidence that HFTs trigger EPMs to benefit from this profitability.

3.1. A typical EPM

To measure HFT activity during EPMs we use directional trade imbalances computed as the difference between trading activity in the direction of the EPM and trading activity in the opposite direction: $HFT^{D} = HFT^{D+} - HFT^{D-}$ and $HFT^{S} = HFT^{S+} - HFT^{S-}$, where HFT^{D} is HFT liquidity demand, HFT^{S} is HFT liquidity supply, and the superscripts + (-) indicate activity in the same (opposite) direction of the EPM return. For example, if HFTs demand 20 shares of liquidity in the direction of the price movement and demand 1 share in the opposite direction, HFT^{D} is +19. Similarly, if HFTs supply 20 shares of liquidity against

[2] Aitken, Cumming and Zhan (2015) find that proliferation of HFT has reduced instances of end-of-day price manipulation.

the direction of the EPM and supply 4 shares in the direction of the EPM, HFT^{S} is -16. We compute similar metrics for nHFTs.

In addition, we introduce two net imbalance metrics, HFT^{NET} ($nHFT^{NET}$) computed as the sum of HFT^{D} and HFT^{S} ($nHFT^{D}$ and $nHFT^{S}$). Since liquidity is typically provided against the direction of return, $(n)HFT^{S}$ usually has a negative value, and the sum of $(n)HFT^{D}$ and $(n)HFT^{S}$ is in effect the difference between liquidity demanding and liquidity providing volume. Net imbalances indicate the direction in which net trading activity by a particular trader type is occurring relative to the EPM direction. For example, a positive (negative) net HFT imbalance indicates overall trading in the direction (opposite) of the EPM.

We begin the discussion of HFT and nHFT activity around EPMs with an illustration. Figure 3 reports the cumulative return (CRET) as well as HFT^{D}, $nHFT^{D}$ and HFT^{NET} starting 100 seconds prior to an average EPM and up to 100 seconds afterwards. We make the following expositional choices. First, the figure includes both positive and negative EPMs, and we invert the statistics for the latter. Second, we benchmark the signs for HFT and nHFT activity to the EPM return. For example, if the EPM return is positive, a negative HFT^{D} ten seconds after the EPM, as in Figure 3, means that HFTs sell the stock via liquidity demanding orders, effectively counteracting the effects of the positive EPM that occurred ten seconds earlier.

INSERT FIGURE 3 ABOUT HERE

Figure 3 shows that prices are generally flat prior to an EPM, then change significantly during the EPM interval, and then revert somewhat during the remaining 100 seconds (10 intervals). There is a large increase in $nHFT^{D}$ during the EPM, with a share imbalance of more than 5,500. In the meantime, HFT^{D} is a little over 2,000 shares. More importantly, HFT^{NET} is negative, indicating that HFT liquidity supply offsets HFT liquidity demand and that HFTs absorb volume imbalances created by nHFTs.[3]

The results in Figure 3 provide first evidence on HFT and nHFT behavior around EPMs. In Table 2, we examine EPM event windows in more detail. Specifically, we focus on event windows that span 20 seconds before and after the EPM interval and report liquidity demand and supply statistics for HFTs and nHFTs. We find that HFT^{NET} is statistically significant in the direction opposite of return during interval t (the EPM interval) and the two following intervals. Further, upon splitting HFT activity into demand and supply, we observe that HFTs trade in the direction of the EPM with their liquidity demanding trades (HFT^{D} is 2,215 shares) and in the opposite direction with their liquidity supplying trades (HFT^{S} is 2,515 shares). HFTs provide about 300 shares of net liquidity against the direction of an average EPM. This finding is contrary to the belief held by some market observers that HFTs trade large amounts in the direction of EPMs.

Is 300 shares a quantity that is too small to claim that HFTs stabilize prices? The results in Table 2 are simple averages and therefore do not suggest that HFT liquidity provision is limited to 300 shares per EPM. Rather, 300 is the number of shares that liquidity demanding nHFTs require during an average EPM. A look at the distribution of HFT^{NET} (untabulated) suggests that HFT liquidity provision varies substantially, reaching tens of thousands of shares for some EPMs.

INSERT TABLE 2 ABOUT HERE

[3] The net imbalance metrics are designed so that HFT^{NET}=-$nHFT^{NET}$.

Beyond being liquidity providers during EPMs, do HFTs trigger EPMs? In the 20 seconds prior to an EPM (t-20), HFT and nHFT trades do not show any directionality. However, in t-10 HFTs trade in the direction of the future EPM return and demand 46 shares more than they supply.[4] It appears that HFTs may play a role in triggering EPMs. We examine this possibility in more detail in a subsequent section.

Following the EPM, HFTs continue to trade in the opposite direction of the EPM return, but unlike in interval t they primarily use liquidity demanding trades. Specifically, HFTs demand a net of 122.5 shares and 42.7 shares against the direction of the preceding EPM return in intervals t+10 and t+20. From Figure 3 we know that on average the EPM return reverses in intervals t+10 and t+20 and so HFTs appear to speed up the reversal.

3.2. EPM types: reversals and permanent price changes

Large price movements can be triggered by at least two types of events: information arrival and trade imbalances. A news arrival, for instance, often results in prices adjusting rapidly to incorporate new information. In an efficient market such price movements will be permanent. Alternatively, trade imbalances usually arise because impatient traders submit large volumes of buy or sell orders and push prices away from the fundamental values. Price movements arising from such pressures are transitory and are followed by reversals. Figure 4 presents an illustration.

[4] In Table 2, as in Figure 3, we benchmark the signs of HFT and nHFT volume to the EPM return.

INSERT FIGURE 4 ABOUT HERE

Do HFTs provide liquidity to both EPM types? To answer this question, we divide the sample into transitory and permanent EPMs. The former are characterized by significant, yet temporary, price changes followed by reversals. We identify these as EPMs that reverse by more than 2/3 by the end of the trading day. The latter, permanent EPMs, do not reverse by more than 1/3 by the end of the day. To allow for a clean separation of the two EPM types we exclude the EPMs that reverse by more than 1/3 and less than 2/3. This reduces the number of EPMs by 2.7%. The results are unchanged when we include the omitted EPMs. The results are also robust to using alternative reversal thresholds and alternative intraday time periods.

In Table 3, we examine the characteristics of the two EPM types and HFT activity around them. Despite a significant difference in post-EPM price patterns, other EPM characteristics (i.e., returns, trading activity, HFT participation and spreads) are similar across the two types (Panel A). For instance, the average absolute return is 0.486% during a typical transitory EPM and 0.481% during a permanent EPM.

In Panel B, we describe HFT activity around the two EPM types. Consistent with the full sample results, HFTs provide liquidity to both types during interval t. Notably, ten seconds prior to permanent EPMs HFTs demand liquidity in the direction of the subsequent price movement. Such activity is potentially profitable, as HFTs aggressively buy low (sell high) prior to a positive (negative) EPM and then provide liquidity to buyers (sellers), who arrive during the interval t by selling high (buying low).

INSERT TABLE 3 ABOUT HERE

How do HFTs foresee permanent EPMs? Recent literature suggests that some traders are better at processing pre-announcement information and also may receive leakages or early releases of such information. For instance, Hu, Pan and Wang (2015) show that some traders obtain pre-releases of the Consumer Sentiment Index and trade on this information before the rest of the market. Similarly, Bernile, Hu and Tang (2016) and Kurov, Sancetta, Strasser and Wolfe (2016) find that some traders obtain information about the upcoming news announcements either via leakages or due to superior ability to analyze public information. The data do not allow us to test these suggestions directly. The results point to the possibility that HFTs trade on price-relevant information prior to permanent EPMs.

3.3. EPM magnitude

Although the EPMs in the sample represent the 99.9th percentile of all price movements, the setup may obscure the picture for the largest EPMs, during which HFT activity may differ from what has been discussed so far. Kirilenko, Kyle, Samadi and Tuzun (2015) show that when prices reached extraordinary lows during the 2010 flash crash, HFTs withdrew from liquidity provision. So far, the results suggest that EPMs are not accompanied by similar withdrawals. But what about the largest EPMs? In Table 4, we ask if HFT liquidity provision varies in EPM magnitude, and particularly if HFTs provide liquidity to the largest of the extreme price movements.

Table 4 reports summary statistics and HFT^{NET} results for EPMs divided into four magnitude quartiles, from the relatively small (Q1) to the largest (Q4). As expected trading volume and spreads increase in return magnitude (Panel A). HFT liquidity provision also steadily increases, going from 111 shares in Q1 to 656 shares in Q4 (Panel B). The largest EPMs attract the most HFT liquidity provision. Insofar as these results may be applied to an event like the 2010 Flash Crash, they suggest that it was probably not the magnitude of the crash that triggered HFT withdrawal.

INSERT TABLE 4 ABOUT HERE

3.4. EPM types: standalone and co-EPMs

The 2010 Flash Crash was characterized not only by the large magnitude of price movements but also by the large number of stocks that were affected. It is suggested that liquidity withdrawals during the crash were due to the HFT firms' capital constraints and to their reduced ability to hedge positions in correlated stocks. Capital constraints may restrict HFTs' ability to accumulate large portfolios via market making, whereas the reduced ability to hedge in correlated stocks may increase their risk. The Flash Crash was a uniquely large and rare event and it is not clear if it should be viewed as suggestive of HFT behavior in all instances of multi-stock price movements. To further examine this issue we define co-EPMs as those that occur in two or more stocks during the same 10-second time interval and repeat the prior analyses.

Panel A of Table 5 reports that the sample consists of 43% standalone EPMs and 57% co-EPMs. The prevalence of co-EPMs should not be surprising given the exceptionally high EPM occurrence during the 2008 financial crisis when prices of multiple assets experienced large simultaneous movements (Figure 2). An average co-EPM includes 3.5 stocks. The average return is 0.491% during a standalone EPM and 0.479% during a co-EPM. Trading activity metrics are noticeably different between the two types, with dollar volume during the standalone EPMs being about 75% higher than that during the co-EPMs. The relative spreads are also somewhat higher during the standalone EPMs; 0.085 bps vs. 0.076 bps for the co-EPMs.

Panel B shows that although standalone EPMs are of relatively larger magnitude HFTs provide 1,297 shares of net liquidity during these events. They demand 446 shares of net liquidity during the co-EPMs. These results are consistent with the earlier suggestion that capital constraints and reduced hedging opportunities may prevent HFTs from effectively performing the market making role during co-EPMs even though co-EPMs are of relatively low magnitude.

INSERT TABLE 5 ABOUT HERE

3.5. Does HFT activity during EPMs differ from their usual behavior?

Research shows that HFTs usually demand liquidity in the direction of returns (e.g., Brogaard, Hendershott and Riordan, 2014). If this pattern persisted during EPMs, we would observe significantly positive and large HFT^{NET}. On the contrary, we find that the pattern reverses for an average EPM. Although the pattern does not reverse for co-EPMs it is possible that the positive HFT-return relation is reduced even for these EPMs. Accounting for return magnitude HFTs may demand less liquidity during the times when multiple stocks undergo EPMs than they normally would. To examine this issue we turn to the following multivariate setting:

$$HFT^{NET}{}_{it} = \alpha + \beta_1 1_{EMP_{it}} + \beta_2 Ret_{it} + \beta_3 Vol_{it} + \beta_4 Spr_{it} + \boldsymbol{Lags_{kit-\sigma}\gamma_{k\sigma}} + \varepsilon_{it}, \quad (1)$$

where HFT^{NET} is the difference between HFT^D and HFT^S as discussed earlier; $1_{EPM\,it}$ is a dummy variable equal to one if the 10-second interval t in stock i is identified as an EPM and is equal to zero otherwise, Ret_{it} is the absolute return, Vol_{it} is the traded share volume, Spr_{it} is the percentage quoted spread, and $\boldsymbol{Lags_{kit-\sigma}}$ is a vector of lags for the dependent and each of the independent variables, with $\sigma \in \{1, 2, \ldots, 10\}$. The variables in the vector are indexed with a subscript k. All variables are standardized at the stock level.

Because the coefficients on the 1_{EPM} dummy are related to returns they should be interpreted jointly with those on the Ret variable. For example in column 1 of Table 6, the estimated coefficient on the Ret variable confirms that HFTs usually demand liquidity in the direction of return. In the meantime, the 1_{EPM} dummy shows that HFTs reduce liquidity demand during EPMs, with the incremental effect of -0.818 standard deviations. To understand the economic effect, recall that the variables in eq. 1 are standardized. Also recall from Table 1 that an EPM return is about 10 standard deviations away from the average return. It follows that the total effect of HFT during EPMs is about $0.072 \times 10 - 0.818 =$ -0.098. Consistent with the univariate results HFTs provide net liquidity during an average EPM.

INSERT TABLE 6 ABOUT HERE

Having established the basic result we next turn to HFT activity during the previously identified EPM types. Column 2 shows that during both transitory and permanent EPMs the normally positive HFT-return relation is significantly reduced. In column 3 we find the same result for standalone and co-EPMs, yet the decline is much greater for the standalone EPMs. Similar results emerge in column 4 that accounts for EPM magnitude; the normally positive relation between HFT behavior and returns is reduced, more so during the largest EPMs. Overall, even in cases when they demand liquidity during the EPM episodes (the co-EPM case), HFTs demand considerably less than they normally would.

3.6. HFT-return relation within the 10-second intervals

The 10-second event windows are quite long given the speed of modern trading and may conceal nefarious aspects of HFT behavior. Yang and Zhu (2015) propose and van Kervel and Menkveld (2015) show that HFTs are able to recognize trading patterns after a period of time and switch from supplying liquidity to demanding it. Although van Kervel and Menkveld (2015) focus on time horizons that are much longer than ours, even one second is a long enough time for HFT algorithms to re-evaluate a trading strategy. It is possible that HFTs supply liquidity at the beginning of EPMs yet exacerbate their tail ends.

To examine this possibility in Figure 5 we plot second by second cumulative returns, HFT, and nHFT activity centered on the largest one-second return during an average EPM. The figure shows that prices continue to move in the direction of the largest return for several seconds afterwards. If HFT algorithms had been designed to quickly switch from liquidity supply to demand after observing large price changes they would have had sufficient time to do so. The figure contains no evidence of HFT^{NET} switching to positive values. If anything, it remains slightly negative.

INSERT FIGURE 5 ABOUT HERE

3.7. Profitability of liquidity provision during EPMs

The data show that HFTs usually provide liquidity to nHFTs during both transitory and permanent EPMs. Since HFTs choose to do so liquidity provision should be profitable. How are these profits derived? We provide an example for permanent EPMs when the price moves up. The same logic, but in reverse, applies to negative permanent EPMs. During positive permanent EPMs as described in Figure 4, if a trader limits liquidity provision to the size of his existing long inventory, he will have bought low and sold high. If he provides liquidity indiscriminately, in the amount larger than the existing inventory, he may accumulate a money-losing short position. Table 3 shows that HFTs accumulate some inventory in interval t-10 prior to permanent EPMs. This inventory is smaller than the amount of liquidity they provide in period t. If their inventory before t-10 is insufficient HFTs may lose money during permanent EPMs.

During transitory price movements when the price first moves up and then down (Figure 4) a skilled trader may profit by initially selling high to the impatient buyers and then buying low when the price reverses. The literature shows that providing liquidity during such reversals is profitable (Hendershott and Seasholes, 2007; Nagel, 2012; So and Wang, 2014). This strategy does not require pre-existing inventory as profits are derived from the inventory accumulated during the EPM itself. In summary, it is possible that HFTs profit from both permanent and transitory EPMs. In Table 7 we examine the data for evidence of such profits.

We estimate HFT trading revenues on EPM days and compare them to the days without EPMs. We follow the approach used by Sofianos (1995), Menkveld (2013), and Brogaard, Hendershott and Riordan (2014) and assume that for each sample stock and each day HFTs start and end the day with zero inventory and that all inventory accumulated by the end of the day is sold at the closing midpoint. We compute the revenue from HFT for each stock and each day as:

$$\pi HFT = -\sum_{n=1}^{N} HFT_n \times I_n \times P_n + invHFT_N \times P_N, \tag{2}$$

where HFT_n is the number of shares traded by HFTs during the n^{th} transaction, I is the indicator equal to 1 for buy trades and -1 for sell trades, P_n is the trade price, $invHFT_N$ is the inventory accumulated through HFT trades by the end of the day, and P_N is the end of day midquote. Following Brogaard, Hendershott and Riordan (2014) we adjust transaction prices by the taker fee of \$0.00295 and the maker rebate of \$0.0028, although the results are robust to other levels of maker-taker fees and to omitting the fees. The first term of eq. 2 represents cash flows throughout the day and the second term assigns a value to the end-of-day inventory.

To assess the impact of EPMs on daily HFT revenues we estimate the following panel regression for each stock i on day t:

$$\pi HFT_{it} = \alpha + \beta_1 nTransitory_{it} + \beta_2 nPermanent_{it} + \varepsilon_{it}, \tag{3}$$

where $nTransitory$ and $nPermanent$ are count variables that capture the number of EPMs of each sub-type. An additional specification replaces the count variables for the sub-types with a single count variable $nEPM$, for all EPMs.

The results are reported in Table 7. The intercept in Panel A shows that the average HFT revenue on days without EPMs is $3,672. The average revenue is higher by $3,873 on days with transitory EPMs. The revenue is lower by $3,004 on days with permanent EPMs. Even on permanent EPM days the total revenue is still a positive $668 (=$3,672-3,004). Even though the losses to permanent EPMs are substantial Panel B shows that the incremental revenue from providing liquidity to an average EPM is a positive $274.[5]

INSERT TABLE 7 ABOUT HERE

The results in Table 7 suggest that even though HFTs may accumulate potentially profitable inventory positions prior to permanent EPMs, they end up providing liquidity in amounts larger than this inventory, which negatively affects their profits. In the meantime, on days with transitory EPMs, HFTs make extra profits that are large enough to compensate for the losses incurred during the permanent EPMs. The strategy of providing liquidity to all EPMs indiscriminately appears profitable.

3.8. HFT activity and future EPMs

Results in the earlier sections show that HFT^{NET} is positive prior to permanent EPMs. Is HFT activity sufficiently large to trigger these EPMs? To evaluate this question we use probit regressions to model the probability of an EPM as a function of lagged values of HFT^{NET}, return, volume and spread:

[5] The intercepts in Panels A and B are somewhat different due to a small difference in the samples sizes as discussed in section 3.2.

$$Prob\ (EPM = 1)_{it} = \alpha + \beta_1 HFT_{it-1}^{NET} + \beta_2 Ret_{it-1} + \beta_3 Vol_{it-1} + \beta_4 Spr_{it-1} + \varepsilon_{it}, \quad (4)$$

where all variables are as previously defined and are lagged by one interval.

The results are in Table 8 and show no evidence of HFT being associated with a higher probability of future EPMs. On the contrary HFT is often associated with a lower EPM probability. In column 1 the marginal effect of the HFT^{NET} variable implies that the probability of an EPM decreases by 0.8% of the unconditional probability with every standard deviation increase in pre-EPM HFT^{NET}. Further, the HFT^{NET} coefficient in column 4, the specification for permanent EPMs, suggests that there is no effect on the EPM likelihood from HFT behavior. Recall that the univariate results show non-trivial positive HFT^{NET} values for permanent EPMs at t-10 consistent with HFTs possibly triggering EPMs. The panel probit analysis shows that once the HFTs' usual relation to future returns is taken into account there is no evidence of HFTs' triggering EPMs. HFT behavior prior to EPMs is not different from their usual behavior. Even though HFTs occasionally demand liquidity prior to EPMs their demand is not a sufficient trigger for the EPMs.

INSERT TABLE 8 ABOUT HERE

4. Robustness

In this section, we check the robustness of main results by examining the sensitivity of HFT liquidity provision to the number of stocks undergoing simultaneous EPMs and to the number of EPMs in the same stock during the day. We also review how HFTs behave when returns are not extreme. Finally, we examine EPMs using a volume clock, which allows us to study EPMs in medium and small stocks.

4.1. Number of stocks in a co-EPM

Earlier results show that HFTs provide substantial liquidity to standalone EPMs, yet demand liquidity during co-EPMs. In Panel A of Table 9, we examine HFT sensitivity to the number of stocks undergoing a co-EPM. The sensitivity is rather high; HFT^{NET} goes from -1,297 shares for a standalone EPM to zero for the two-stock co-EPMs and continues to increase to hundreds of shares for co-EPMs that involve three and more stocks. The results therefore suggest that the threshold for HFT liquidity becoming fragile is rather low.

INSERT TABLE 9 ABOUT HERE

4.2. EPM sequences

Given that HFTs do not provide liquidity to co-EPMs, it is possible that they also remain on the sidelines on days with long sequences of EPMs, especially if these EPMs have the same return direction. In Panel B of Table 9, we ask if HFT liquidity provision is sensitive to the number of same-directional EPMs on a given day. The data show that HFTs usually provide liquidity to the first four EPMs in the sequence and reduce liquidity provision if the sequence continues. There is some evidence of positive HFT^{NET} for very long sequences.

4.3. Returns that are not extreme

Regression results show that HFTs usually trade in the direction of contemporaneous returns, yet this relation reverses for an average EPM. It is not clear if this reversal happens exactly at the chosen 99.9% threshold. In Panel C of Table 9 we report HFT^{NET} for several groups of returns: zero returns, four non-zero return quartiles up to the 99.0^{th} percentile, four groups between the 99.0^{th} and 99.9^{th} percentiles, and finally for the 99.9^{th} percentile. The results show that the positive HFT-return relation exists for virtually all return magnitudes below the 99.9^{th} percentile. Liquidity provision of the type that we document is limited to the largest price movements.

4.4. Volume buckets

Easley, López de Prado and O'Hara (2012, 2016) suggest that additional insights may be gained if modern markets are examined using volume bucketing, which involves parsing data into intervals of equal trading volume. They argue that using time intervals may oversample events like EPMs because volatility tends to cluster. The results from the Lee and Mykland (2012) procedure that we report in the appendix is one way to deal with volatility clustering. In this section we incorporate the role of volatility clustering by examining volume buckets. In addition to mitigating volatility clustering volume bucketing allows us evaluate extreme price movements in medium and small stocks. Recall that the time-based intervals usually do not contain enough trades for EPM identification in such stocks.

To identify extreme price movements using volume bucketing we first estimate the average daily volume for each stock during the two-year sample period. Next we select volume buckets that represent $1/100^{th}$ of the average daily volume. On most days there is residual end-day volume that is not large enough to form its own bucket. To preserve information in this volume we merge it with the last complete bucket of the day. When a single trade has volume that is larger than the bucket size we use this trade to form a single bucket. The resulting volume buckets contain an average of 269 trades in large stocks, 22 in medium and 9 in small stocks.

Finally, we calculate transaction-based absolute returns for each bucket and use the 99.9th percentile cut-off to identify volume-based EPMs (vEPMs).[6]

Table 10 reports the descriptive statistics and the HFT^{NET} results computed during vEPMs for large, medium and small stocks. vEPMs are larger than the original EPMs, with the average return of 1.52% in large stocks, more than 3 times the return reported for the main sample. Medium and small stock returns are 3.44% and 7.00%, respectively. vEPMs are also longer than the 10-second EPMs in the main sample, with lengths varying from 2.5 to about 20 minutes. The results show that vEPMs and EPMs are different types of extreme price movements with EPMs more reflective of the fast-paced trading in modern markets. Nonetheless, it is still interesting to learn how HFTs behave during the slower, larger vEPMs.

INSERT TABLE 10 ABOUT HERE

Du and Zhu (2015) show theoretically that when liquidity events develop relatively quickly HFT participation in liquidity provision is higher because other traders arrive at the market relatively slowly. For slow-paced events the share of HFT participation in liquidity provision should be smaller. The results corroborate these predictions as the data contain only weak evidence of HFTs' acting as net liquidity providers during vEPMs (HFT^{NET} is negative, yet statistically insignificant in stocks of all sizes). Even during these events HFTs do not act as net liquidity demanders. In unreported probit results we also find no evidence of HFTs triggering vEPMs.

[6] The results are the same for larger volume buckets that represent 1/75th and 1/50th of the average daily volume. A finer bucket definition (e.g., 1/500th or 1/1000th of daily volume) does not work for the medium and small stocks, because such buckets do not contain a sufficient number of trades.

5. Conclusion

We provide novel evidence on the stability of liquidity supply by high frequency traders (HFTs), a dominant subset of liquidity providers in modern markets. HFTs are endogenous liquidity providers (ELPs) and do not have the obligation to supply liquidity during stressful times. We show that HFTs are net suppliers of liquidity to non-HFTs (nHFTs) during extreme price movements (EPMs). HFTs provide liquidity to impatient nHFTs, whose trade imbalances cause prices to move rapidly and substantially. HFTs supply liquidity to EPMs, including the most extreme ones and the ones that result in permanent price changes. However, HFT liquidity supply is fragile as they often take liquidity when multiple stocks undergo simultaneous EPMs.

HFTs earn positive revenues on days with EPMs, mostly through the contrarian channel, whereby they supply liquidity to impatient traders by selling high (buying low) and then covering positions by buying low (selling high). Despite their ability to profit from an average EPM the results show that HFTs do not appear to cause EPMs.

Theory suggests that traders like HFTs may choose several ways of reacting to trade imbalances. Traders described by Grossman and Miller (1988) supply liquidity during trade imbalances and benefit from price reversals that follow. The predatory traders of Brunnermeier and Pedersen (2005) demand liquidity alongside trade imbalances. The back-runners of Yang and Zhu (2015) supply liquidity until they recognize an institutional trading pattern and then switch to demand liquidity. In our setting HFT behavior is more consistent with that described by Grossman and Miller (1988), although the data point to occasional back-running for long EPM sequences.

References

Aitken, Cumming and Zhan (2015). High frequency trading and end-of-day price dislocation. Journal of Banking and Finance 59, 330-349.

Anand, A., and K. Venkataraman (2015). Market conditions, fragility, and the economics of market making. Journal of Financial Economics, forthcoming.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens (2001). The distribution of realized stock return volatility. Journal of Financial Economics 61, 43–76.

Angel, J., L. Harris, and C. Spatt (2011). Equity trading in the 21st century. Quarterly Journal of Finance 1, 1-53.

Bernile, G., J. Hu, and Y. Tang (2016). Can information be locked up? Informed trading ahead of marco-news announcements. Journal of Financial Economics, forthcoming.

Bessembinder, H., J. Hao, and M. Lemmon (2011). Why designate market makers? Affirmative obligations and market quality. Working paper.

Bongaerts, D. and M. Van Achter (2015). High-frequency trading and market stability. Working paper.

Brogaard, J. A., T. Hendershott, and R. Riordan (2014). High-frequency trading and price discovery. Review of Financial Studies 27, 2267-2306.

Brunnermeier, M and L. H. Pedersen (2005). Predatory trading. Journal of Finance 60, 1825-1863.

Carrion, A. (2013). Very fast money: High-frequency trading on the NASDAQ. Journal of Financial Markets 16, 680-711.

Cespa, J. and X. Vives (2015). The welfare impact of high frequency trading. Working paper.

Chaboud, A., B. Chiquoine, E. Hjalmarsson, and C. Vega (2014). Rise of machines: Algorithmic trading in the foreign exchange market. Journal of Finance 2045-2084.

Chan, K. C., W. G. Christie, and P. H. Schultz (1995). Market structure and the intraday pattern of bid-ask spreads for NASDAQ securities. Journal of Business 68, 35-60.

Chordia, T., A. Goyal, B. N. Lehmann, and G. Saar (2013). High-frequency trading. Journal of Financial Markets 16, 637-645.

Clark-Joseph, A., M. Ye and C. Zi (2016). Designated market makers still matter: Evidence from two natural experiments. Working paper.

Colliard, J.-E. (2015). Catching falling knives: Speculating on market overreaction. Working paper.

Conrad, J., S. Wahal, and J. Xiang (2015). High-frequency quoting, trading, and the efficiency of prices. Journal of Financial Economics 116, 271-291.

Duffie, D., and J. Pan (2001). Analytical value-at-risk with jumps and credit risk. Finance and Stochastics 5, 155-180.

Du, S., and H. Zhu (2015). Welfare and optimal trading frequency in dynamic double auctions. Working paper.

Easley, D., M. López de Prado, and M. O'Hara (2012). Flow toxicity and liquidity in a high-frequency world. Review of Financial Studies 25, 1457-1493.

Easley, D., de Prado, M. L., & O'Hara, M. (2016). Discerning information from trade data. Journal of Financial Economics 120, 269-285.

Egginton, J. (2014). The declining role of NASDAQ market makers. Financial Review 49, 461-480.

Eraker, B., M. Johannes, and N. Polson (2003). The impact of jumps in volatility and returns. Journal of Finance 58, 1269-1300.

Golub, A., J. Keane and S-H. Poon (2013). High-frequency trading and mini flash crashes. Working paper.

Gromb, D., and D. Vayanos (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. Journal of Financial Economics 66, 361-407.

Grossman, S. and M. Miller (1988). Liquidity and market structure. Journal of Finance 43, 617-633.

Harris, L. (2013). What to do about high-frequency trading. Financial Analysts Journal 69, 6-9.

Hasbrouck, J., and G. Saar (2013). Low-latency trading. Journal of Financial Markets 16, 646-679.

Hendershott, T. and M. Seasholes (2007). Market maker inventories and stock prices. American Economic Review 97, 210-214.

Holden, C., and S. Jacobsen (2014). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. Journal of Finance 69, 1747-1785.

Hu, G. X., J. Pan, and J. Wang (2015). Early peek advantage?. Working paper.

Jarrow, R., and E. Rosenfeld (1984). Jump risks and the intertemporal capital asset pricing model. Journal of Business 3 (57), 337-351.

Jones, C. (2013). What do we know about high-frequency trading?. Working paper.

Kirilenko, A., A. S. Kyle, M. Samadi, and T. Tuzun (2015). The flash crash: The impact of high-frequency trading on an electronic market. Working paper.

Korajczyk, R., and D. Murphy (2015). High frequency market making to large institutional trades. Working paper.

Kurov, A., A. Sancetta, G. Strasser, and M. H. Wolfe (2016). Price drift before U.S. macroeconomic news: Private information about public announcements?. Working paper.

Leal, S. J., M. Napoletano, A. Roventini, and G. Fagiolo (2014). Rock around the clock: An agent-based model of low- and high-frequency trading. Working paper.

Lee, S. S., and P. A. Mykland (2012). Jumps in equilibrium prices and market microstructure noise. Journal of Econometrics 168, 396–406.

Liu, J., F. Longstaff, and J. Pan (2003). Dynamic asset allocation with event risk. Journal of Finance 58, 231-259.

Malinova, K., A. Park, and R. Riordan (2014). Do retail traders suffer from high-frequency traders?. Working paper.

Menkveld, A. (2013). High-frequency trading and the new-market makers. Journal of Financial Markets 16, 712-740.

Menkveld, A., and B. Yueshen (2015). The flash crash: A cautionary tale about highly fragmented markets. Working paper.

Nagel, S. (2012). Evaporating liquidity. Review of Financial Studies 25, 2005-2039.

O'Hara, M., C. Yao, and M. Ye (2014). What's not there: Odd-lots and market data. Journal of Finance 69, 2199-2236.

Raman, V., M. Robe, and P. Yadav (2014). Electronic market makers, trader anonymity and market fragility. Working paper.

So, E., and S. Wang (2014). News-driven return reversals: Liquidity provision ahead of earnings announcements. Journal of Financial Economics 114, 20-35.

Sofianos, G., 1995, Specialist gross trading revenues at the New York Stock Exchange. Working paper.

van Kervel, V., and A. Menkveld (2015). High-frequency trading around large institutional orders. Working paper.

Yang, L., and H. Zhu (2015). Back-running: Seeking and hiding fundamental information in order flows. Working paper.

Table 1. Summary statistics

The table reports summary statistics for the sample of extreme price movements (EPMs) (Panel A) and for the full sample of 10-second intervals (Panel B). $Absolute\ Return$ is the absolute value of the 10-second midpoint return. $Total\ (HFT)\ Trades$ is the number of (HFT) trades during the interval. $Dollar\ Volume$ and $Share\ \ Volume$ are the total dollar and share volume traded during the interval. $Quoted\ Spread$ and $Relative\ Spread$ are quoted and relative quoted NBBO spreads, respectively in dollars and percentage points. All statistics are averaged over the 10-second sampling intervals.

Panel A: Extreme price movements

	Mean	**Median**	**Std. Dev.**
Absolute Return, %	0.484	0.441	0.193
Total Trades	73.0	43.0	88.7
Total HFT Trades	57.6	33.0	73.2
Dollar Volume	473,232	171,158	1,024,504
Share Volume	15,595	5,431	31,734
Quoted Spread, $	0.046	0.016	0.147
Relative Spread, %	0.080	0.065	0.148
N	45,406		

Panel B: Full sample

Absolute Return, %	0.028	0.009	0.048
Total Trades	18.1	11.0	18.7
Total HFT Trades	15.8	10.0	15.5
Dollar Volume	76,285	14,038	231,397
Share Volume	1,991	318	6,055
Quoted Spread, $	0.026	0.010	0.057
Relative Spread, %	0.046	0.040	0.033
N	45.4 M		

Table 2. Liquidity supply and demand around EPMs

The table reports directional trading volume around extreme price movements. Time interval t is the 10-second EPM interval. In addition, we report the results for the two time intervals preceding the EPM and two subsequent time intervals. HFT^{D} ($nHFT^{D}$) is the difference in liquidity-demanding HFT (nHFT) volume in the direction of the EPM and liquidity-demanding volume against the direction of the EPM. HFT^{S} ($nHFT^{S}$) is the difference in liquidity-providing volume against the direction of the EPM and liquidity-providing volume in the direction of the EPM. HFT^{NET} ($nHFT^{NET}$) is the difference between HFT^{D} and HFT^{S} ($nHFT^{D}$ and $nHFT^{S}$). *p*-values are in parentheses. *** and ** indicate statistical significance at the 1% and 5% levels.

	t-20	**t-10**	**t**	**t+10**	**t+20**
HFT^{NET}	1.5	45.7**	-299.3***	-122.5***	-42.7**
	(0.94)	(0.04)	(0.00)	(0.00)	(0.04)
HFT^{D}	30.6	163.4***	2215.2***	-279.0***	-99.1***
	(0.13)	(0.00)	(0.00)	(0.00)	(0.00)
HFT^{S}	-29.1	-117.6***	-2514.6***	156.5***	56.4***
	(0.14)	(0.00)	(0.00)	(0.00)	(0.00)
$nHFT^{NET}$	-1.5	-45.7**	299.3***	122.5***	42.7**
	(0.94)	(0.04)	(0.00)	(0.00)	(0.04)
$nHFT^{D}$	75.3**	326.7***	5576.3***	672.4***	317.0***
	(0.03)	(0.00)	(0.00)	(0.00)	(0.00)
$nHFT^{S}$	-76.8**	-372.5***	-5277.0***	-549.9***	-274.3***
	(0.02)	(0.00)	(0.00)	(0.00)	(0.00)

Table 3. Transitory and permanent EPMs

The table reports summary statistics for transitory and permanent EPMs. Transitory EPMs are those that revert by more than 2/3 of the EPM return by the end of the trading day. Permanent EPMs are those that do not revert by more than 1/3 by the end of the trading day. Because we exclude EPMs that revert by the amount between 1/3 and 2/3, the total number of EPMs in this table is 2.7% lower than the number reported in Panel A of Table 1. Panel B reports HFT^{NET} around the two EPM types.

Panel A: Summary statistics

	transitory		permanent	
	mean	**std. dev.**	**mean**	**std. dev.**
Absolute Return, %	0.486	0.195	0.481	0.191
Total Trades	72.81	89.07	72.11	86.60
Total HFT Trades	57.26	72.83	57.11	72.39
Dollar Volume	472,562	1,052,698	460,269	960,689
Share Volume	15,396	31,448	15,327	30,076
Quoted Spread, $	0.047	0.150	0.046	0.142
Relative Spr., %	0.081	0.146	0.080	0.151
N	21,250		22,913	

Panel B: HFT^{NET}

	t-20	**t-10**	**t**	**t+10**	**t+20**
transitory	-7.0	-30.1	-339.7***	-149.1***	-48.7
permanent	3.9	118.6***	-258.4***	-96.4***	-30.8

Table 4. EPM magnitude quartiles

Panel A divides EPMs into quartiles by return magnitude, from smallest to largest. Panel B contains HFT^{NET} statistics.

Panel A: Summary statistics

	Q1 (small)		Q2	
	mean	**std. dev.**	**mean**	**std. dev.**
Absolute Return, %	0.387	0.094	0.419	0.103
Total Trades	61.31	68.58	64.15	71.25
Total HFT Trades	48.96	57.83	50.96	59.05
Dollar Volume	378,141	798,985	408,766	897,376
Share Volume	12,487	24,759	12,973	24,216
Quoted Spread, $	0.042	0.134	0.043	0.111
Relative Spr., %	0.074	0.086	0.075	0.083
N	11,358		11,327	
	Q3		**Q4 (large)**	
Absolute Return, %	0.471	0.118	0.659	0.268
Total Trades	70.81	80.88	95.75	120.11
Total HFT Trades	55.90	66.96	74.48	98.58
Dollar Volume	452,857	932,231	652,912	1,356,125
Share Volume	15,070	30,330	21,842	43,031
Quoted Spread, $	0.046	0.136	0.055	0.190
Relative Spr., %	0.080	0.131	0.090	0.238
N	11,358		11,363	

Panel B: HFT^{NET}

	t-20	t-10	t	t+10	t+20
Q1	-29.8	-66.5	-110.8*	-125.0***	4.9
Q2	16.4	99.6***	-145.5***	-82.2**	-61.8*
Q3	24.9	66.2	-293.7***	-82.8*	-56.6
Q4	-11.1	82.5	-655.5***	-203.6***	-60.8

Table 5. Standalone and co-EPMs

Panel A divides EPMs into standalone and co-EPMs, with the latter group capturing EPMs that occur simultaneously in several stocks. Panel B contains HFT^{NET} statistics.

Panel A: Summary statistics

	standalone		co-EPMs	
	mean	**std. dev.**	**mean**	**std. dev.**
Absolute Return, %	0.491	0.198	0.479	0.190
Total Trades	89.30	107.05	60.83	69.54
Total HFT Trades	68.60	87.76	49.34	58.72
Dollar Volume	625,553	1,272,083	359,359	770,887
Share Volume	21,368	40,535	11,280	22,092
Quoted Spread, $	0.049	0.125	0.044	0.160
Relative Spr., %	0.085	0.118	0.076	0.168
# stocks			3.5	2.66
N	19,424		25,982	

Panel B: HFT^{NET}

	t-20	t-10	t	t+10	t+20
standalone	-2.2	-32.1	-1296.9***	-128.4***	-40.9
co-EPMs	4.4	103.9***	446.4***	-118.1***	-44.1**

Table 6. Net HFT activity and EPMs

The table reports estimated coefficients from the following regression:

$$HFT^{NET}{}_{it} = \alpha_i + \beta_1 1_{EMP}{}_{it} + \beta_2 Ret_{it} + \beta_3 Vol_{it} + \beta_4 Spr + \boldsymbol{Lags_{kit-\sigma}\gamma_{k\sigma}} + \varepsilon_{it},$$

where HFT^{NET} is the difference between HFT^{D} and HFT^{S}; the dummy 1_{EPM} is equal to one if a 10-second interval t is identified to contain an EPM and is equal to zero otherwise; $1_{EPM\text{-}TRANSITORY}$ and $1_{EPM\text{-}PERMANENT}$ are dummies that capture the two EPM types; $1_{EPM\text{-}STANDALONE}$ captures the standalone EPMs; $1_{CO\text{-}EPM}$ captures EPMs that occur simultaneously in two or more sample stocks; $1_{EPM\text{-}Q1}$ through $1_{EPM\text{-}Q4}$ identify four EPM quartiles by magnitude, from the smallest to the largest; Ret is the absolute return; Vol is the total trading volume; Spr is the percentage quoted spread; and $\boldsymbol{Lags_{kit-\sigma}}$ is a vector of σ lags of the dependent variable and each of the independent variables, with $\sigma \in \{1, 2, \ldots, 10\}$ and the variables indexed with a subscript k. All non-dummy variables are standardized on the stock level. Regressions are estimated with stock fixed effects. p-Values associated with the double-clustered standard errors are in parentheses. *** and ** denote statistical significance at the 1% and 5% levels.

	(1)	(2)	(3)	(4)
1_{EPM}	-0.818***			
	(0.00)			
$1_{EPM\text{-}TRANSITORY}$		-0.861***		
		(0.00)		
$1_{EPM\text{-}PERMANENT}$		-0.776***		
		(0.00)		
$1_{EPM\text{-}STANDALONE}$			-1.441***	
			(0.00)	
$1_{CO\text{-}EPM}$			-0.328***	
			(0.00)	
$1_{EPM\text{-}Q1}$				-0.490***
				(0.00)
$1_{EPM\text{-}Q2}$				-0.631***
				(0.00)
$1_{EPM\text{-}Q3}$				-0.807***
				(0.00)
$1_{EPM\text{-}Q4}$				-1.406***
				(0.00)
Ret	0.072***	0.072***	0.072***	0.073***
	(0.00)	(0.00)	(0.00)	(0.00)
Vol	0.081***	0.081***	0.081***	0.081***
	(0.00)	(0.00)	(0.00)	(0.00)
Spr	-0.010***	-0.010***	-0.010***	-0.010***
	(0.00)	(0.00)	(0.00)	(0.00)
Adj. R^2	0.02	0.02	0.02	0.02

Table 7: HFT revenues on EPM days

Panel A reports coefficient estimates from the following regression model:

$$\pi HFT_{it} = \alpha + \beta_1 nTransitory_{it} + \beta_2 nPermanent_{it} + \varepsilon_{it},$$

where πHFT_{it} is the total revenue from net HFT activity in stock i on day t, and and $nTransitory$ and $nPermanent$ are count variables for the number of each of the two EPM types in stock i day t. Panel B reports coefficient estimates from a similar model that does not differentiate among the EPM types. Profits are computed as follows:

$$\pi HFT = -\sum_{n=1}^{N} HFT_n \times I_n \times P_n + invHFT_N \times P_N,$$

where HFT_n is the number of shares traded by HFTs during the n^{th} transaction, I is the indicator equal to 1 for buy trades and -1 for sell trades, P_n is the trade price, $invHFT_N$ is the inventory accumulated through HFT trades by the end of the day, and P_N is the end of day midquote. The regression is estimated without fixed effects, to maintain a meaningful intercept that captures average HFT^{NET} profits on days without EPMs.

Panel A: Profitability by EPM type

	estimate	**p-value**
Intercept	3672.41	(0.00)
nTransitory	3873.11	(0.00)
nPermanent	-3004.37	(0.00)

Panel B: Overall EPM profitability

	estimate	**p-value**
Intercept	3718.09	(0.00)
nEPM	273.89	(0.01)

Table 8. EPM determinants

The table reports the coefficients and the marginal effects from a probit model of EPM occurrence:

$$Prob\ (EPM = 1)_{it} = \alpha + \beta_1 HFT^{NET}_{it-1} + \beta_2 Ret_{it-1} + \beta_3 Vol_{it-1} + \beta_4 Spr_{it-1} + \varepsilon_{it},$$

where the dependent variable is equal to one if an interval t contains an extreme price movement and zero otherwise. All independent variables are lagged by one interval. HFT^{NET} is the share volume traded in the direction of the price movement minus the share volume traded against the direction of the price movement for all HFT trades, Ret is the absolute return, Vol is total traded volume, Spr is the percentage quoted spread. All variables are standardized on the stock level. The marginal effects are scaled by a factor of 1,000. p-Values are in parentheses. *** and ** indicate statistical significance at the 1% and 5% levels.

	All	**Standalone**	**Co-EPMs**	**Permanent**	**Transitory**
	(1)	(2)	(3)	(4)	(5)
Intercept	-3.232***	-3.438***	-3.380***	-3.403***	-3.423***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
HFT^{NET}_{t-1}	-0.003***	-0.006***	0.001	-0.001	-0.005***
Marginal Effect	-0.008	-0.009	0.001	-0.002	-0.008
	(0.00)	(0.00)	(0.42)	(0.25)	(0.00)
Controls	Yes	Yes	Yes	Yes	Yes
Pseudo-R^2	0.14	0.11	0.13	0.12	0.12

Table 9. Co-EPMs, EPM sequences and the HFT-return relation

The table reports directional HFT^{NET} for standalone and co-EPMs, EPM sequences, and for various return percentiles. Panel A divides EPMs into standalone and co-EPMs, with column (1) showing the number of stocks undergoing an EPM in any particular interval. Panel B examines sequences of same-directional EPMs during the trading day, with column (4) identifying the position of a particular EPM in the sequence. Panel C reports HFT^{NET} for price movements divided into absolute return groups: zero returns, four return quartiles up to 99.0th percentile, four percentile groups between 99.0th and 99.9th percentiles and finally the 99.9th percentile. p-values are in parentheses. ***, ** and * indicate statistical significance at the 1%, 5% and 10% levels.

Panel A: standalone and co-EPMs			**Panel B: EPM sequences**			**Panel C: return percentiles**	
	HFT^{NET}	**# obs.**		**HFT^{NET}**	**# obs.**		**HFT^{NET}**
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1	-1,297***	19,424	1st	-787***	10,447	=0	3***
2	-2	7,476	2nd	-583***	5,797		
3	341***	4,434	3rd	-403***	4,028	(0; 25]	28***
4	340***	3,104	4th	-352***	3,073	(25; 50]	41***
5	497***	2,210	5th	-182	2,429	(50; 75]	80***
6	453***	1,740	6th	-258*	2,014	(75; 99.0]	122***
7	869***	1,211	7th	40	1,730		
8	617***	1,008	8th	44	1,483	(99.0-99.25]	66***
9	669***	792	9th	72	1,297	(99.25-99.5]	106***
10	635***	640	10th	-36	1,128	(99.5-99.75]	30*
11+	1,351***	3,367	11th+	124**	11,980	(99.75-99.9)	-22
						99.9+	-299***

Table 10. Volume buckets

The table reports summary statistics for the samples of large, medium and small stocks during the extreme price movements identified via the volume-clock approach. Ret is the average volume-based EPM (vEPM) return calculated using transaction prices, Length is the vEPM length, #Trades is the number of trades, Vol is the number of shares traded, HFT^{NET} is the mean difference between HFT demand and HFT supply, # obs. is the number of vEPMs. p-values are in parentheses.

	Ret, %	**Length, sec.**	**# Trades**	**Vol, sh.**	**HFT^{NET}**	***p*-value**	**# obs.**
large	1.52	152	269	45,400	-102.7	(0.60)	2,002
medium	3.44	1,228	22	3,201	-2.8	(0.93)	2,067
small	7.00	1,144	9	1,234	-17.7	(0.15)	1,744

Figure 1: Intraday distribution of EPMs

The figure contains a minute-by-minute intraday distribution of EPMs identified using the 99.9 technique.

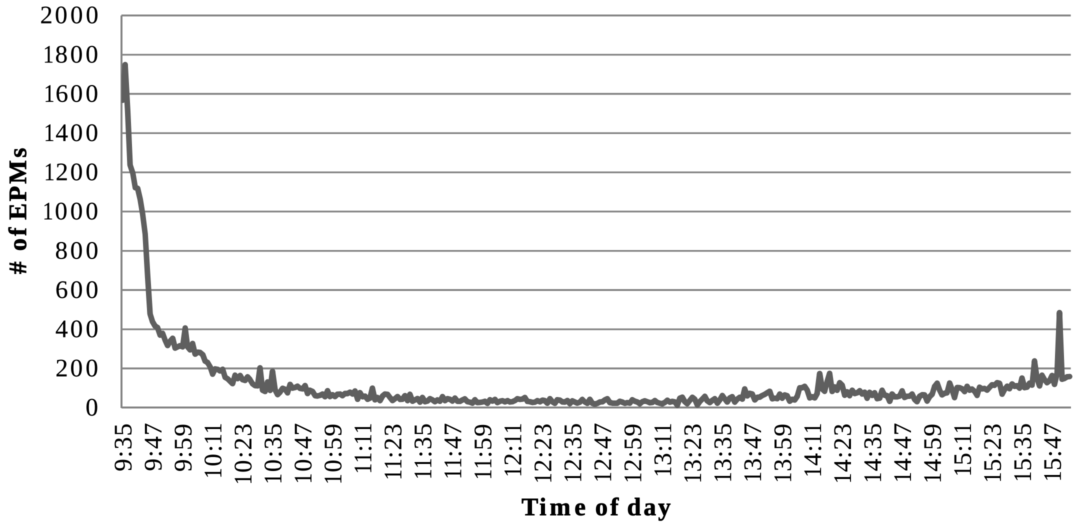

Figure 2: Daily distribution of EPMs

The figure contains the daily distribution of 45,406 sample EPMs identified during the 2008-2009 period using the 99.9 technique.

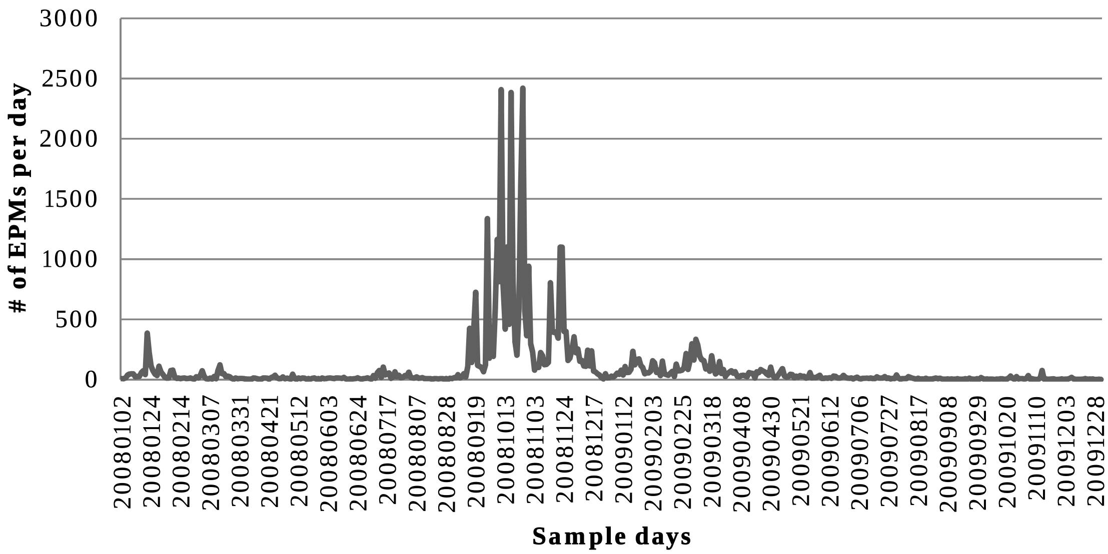

Figure 3: HFT and nHFT activity around EPMs

The figure displays the average return path and trading activity around 45,406 sample EPMs. HFT^D ($nHFT^D$) is liquidity demanded by HFTs (nHFTs) in the direction of the EPM (in # shares) minus liquidity demanded against the direction of the EPM. HFT^{NET} is the net effect of HFT liquidity demand and supply. CRET is the cumulative return. The figure includes both positive and negative EPMs, and for exposition purposes we invert the statistics for the latter.

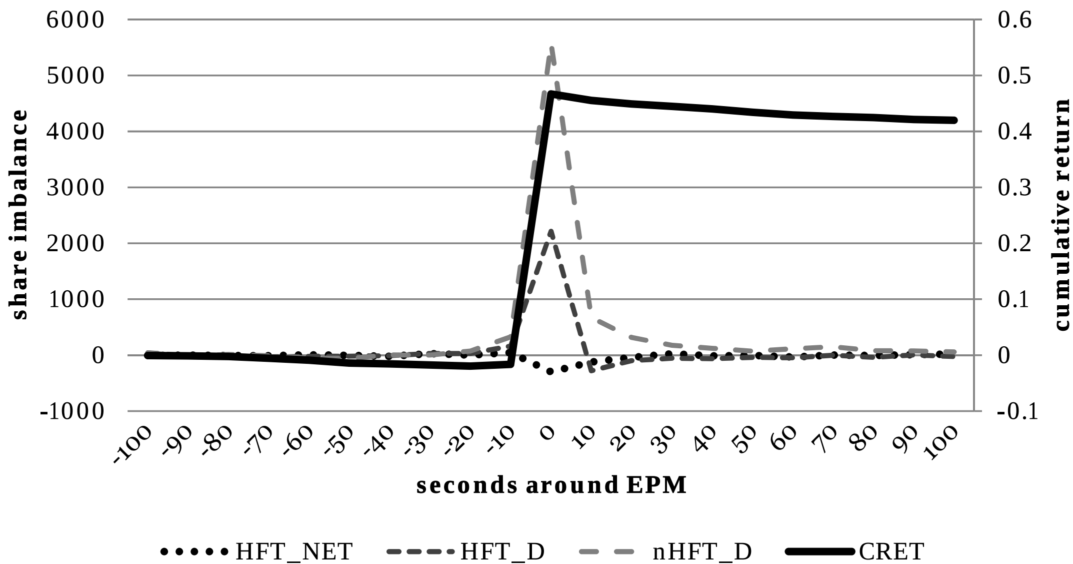

Figure 4: EPM types, an illustration

The figure describes two EPM types according to the associated price patterns: (a) a transitory EPM that reverses by the end of the trading day and (b) a permanent EPM that does not reverse.

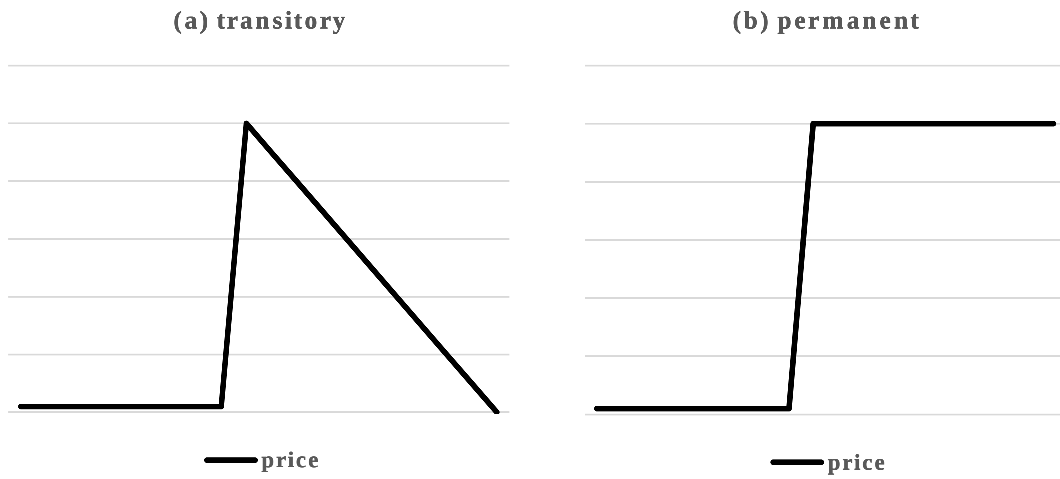

Figure 5: HFT and nHFT activity during EPMs, a second by second view

The figure displays the average second by second price path and trading activity during [-10; +10]-second windows centered on the largest one-second EPM return. HFT^D ($nHFT^D$) is liquidity demanded by HFTs (nHFTs) in the direction of the EPM (in # shares) minus liquidity demanded against the direction of the EPM. HFT^{NET} is the net effect of HFT liquidity demand and supply. CRET is the cumulative return. The figure includes both positive and negative EPMs, and for exposition purposes we invert the statistics for the latter.